QUBAHAN EDUCATIONAL JOURNAL
VOL. 3, NO. 4, October 2026
https://doi.org/10.48161/qej.v3n4a22

Stochastic Decision Support under Uncertainty: Markov Transitions and Monte Carlo Simulation

Maryna Shlenova[1*], Mercy Ifunanya Ani[2]

[1] Ukrainian Language and Language Communication, National Aerospace University "Kharkiv Aviation Institute", Kharkiv, Ukraine;
[1] Department of Science Education, Faculty of Education, Alex Ekwueme Federal University Ndufu-Alike, Ikwo, Ebonyi State, Nigeria;

* **Corresponding author:** science.edu.ua@gmail.com.

ABSTRACT: This study develops and evaluates a stochastic-process modeling framework for decision support under uncertainty. Random processes in operational systems evolve over time, yet many decision-support procedures still rely on deterministic summaries that do not show transition risk, future-state uncertainty, or the probability of adverse outcomes. The proposed framework integrates probability distribution fitting, Markov state-transition modeling, and Monte Carlo simulation to convert uncertain system behavior into measurable decision indicators. We use two empirical cases: hourly traffic-flow data with 48,120 observations from four junctions and a predictive-maintenance dataset with 10,000 machine-operation observations. In the traffic case, we transform vehicle counts into low, moderate, high, and critical congestion states and estimate transition matrices to model next-hour movement between states. In the maintenance case, operating variables are converted into normal, watch, high-risk, and failure states to examine risk escalation. Results show that the lognormal distribution provides the strongest parametric fit for vehicle counts, while the traffic Markov model achieves a one-step state accuracy of 0.720, equal to the persistence baseline but with additional probabilistic interpretability. Monte Carlo simulation indicates a high probability of encountering at least one critical traffic state within 24 hours when the current state is already critical. In the maintenance case, the probabilistic validation model achieves an area under the receiver operating characteristic curve (ROC-AUC) of 0.976 and an F1-score of 0.758. The findings show that stochastic modeling can support risk-sensitive decisions by reporting transition probabilities, uncertainty intervals, and scenario-based future outcomes.

Keywords: stochastic modeling; random processes; decision support; Markov chains; Monte Carlo simulation.

I. INTRODUCTION

Operational systems rarely move deterministically. Traffic volumes change from hour to hour because of demand cycles, local events, weather, and unobserved driver behavior. Engineering assets move from normal operation to warning conditions, and sometimes to failure, due to cumulative wear, unstable loads, environmental effects, and irregular process conditions. In such settings, decision-makers require more than a single prediction. They need to know how likely a system is to remain stable, how likely it is to move into a risky state, and how alternative decisions may perform under uncertain future conditions.

Decision-support systems translate data into usable operational recommendations. However, many applied systems still present averages, deterministic forecasts, or classification labels without showing the uncertainty behind them. A traffic control unit may receive an expected vehicle count but not the probability of a critical congestion state. A maintenance manager may receive a failure-risk label but not the probability

that an asset will enter a higher-risk condition within a defined operating horizon. These limitations reduce the practical value of data-driven decision support, especially when outcomes are dynamic and time-dependent.

Stochastic modeling provides a suitable mathematical foundation for this problem because it treats system behavior as a random process evolving over time. Instead of assuming that future behavior is known exactly, a stochastic framework represents possible future states and attaches probabilities to them. This allows decision-makers to evaluate expected outcomes, risk probabilities, transition patterns, and uncertainty intervals. When combined with simulation, stochastic modeling can also generate future scenarios that support comparison between decision alternatives.

This study addresses the need for a practical, reproducible stochastic decision-support framework. The framework is intentionally built from interpretable components: distribution fitting, state design, Markov transition matrices, Monte Carlo simulation, and risk-based decision indicators. These components are mathematically rigorous but still understandable to analysts and managers who need actionable evidence rather than opaque predictions. The study is also empirical. It uses traffic-flow data as the main time-dependent case and predictive-maintenance data as a validation case to show that the same logic can be transferred from transportation operations to engineering maintenance.

The central problem this study addresses is that random system behavior is often measured but not fully transformed into decision-support information. A dataset may contain thousands of time-indexed observations, but unless it models movement between states, the decision-maker cannot easily understand the probability of risk escalation. Likewise, a failure dataset may support high predictive accuracy, but without stochastic states and scenario simulation, it may not clearly support maintenance timing or risk prioritization. The proposed framework fills this gap by connecting data patterns to state transitions and future scenarios.

The paper contributes to the applied stochastic modeling and decision-support literature in three ways. First, it presents a unified framework that links probability distributions, Markov chains, and Monte Carlo simulation in a single decision-support procedure. Second, it demonstrates the procedure using two different operational domains, which strengthens the generality of the framework. Third, it reports decision-oriented outputs such as transition probabilities, the probability of a critical state, expected future value, and model validation indicators, rather than relying only on point predictions.

1. RESEARCH AIM AND RESEARCH QUESTIONS

The aim of this study is to develop and empirically evaluate a stochastic process modeling framework that supports decision-making under uncertainty by estimating state transitions, simulating future scenarios, and reporting interpretable risk indicators.

The study is guided by four research questions. RQ1 asks how random system behavior can be transformed into stochastic states suitable for decision support. RQ2 asks how Markov transition probabilities describe movement between normal, warning, high-risk, and critical states. RQ3 asks how Monte Carlo simulation can estimate future risk under uncertain system evolution. RQ4 asks whether the same stochastic decision-support framework remains useful across both traffic-flow and predictive-maintenance applications.

II. LITERATURE REVIEW

Probabilistic forecasting provides the conceptual basis for decision support under uncertainty because it treats future outcomes as distributions rather than fixed values. Gneiting and Katzfuss [6] argue that useful forecasting should report calibrated probability distributions, while Kroese et al. [9] show that Monte Carlo simulation has become a central tool in engineering, finance, operations, and reliability analysis. Hüllermeier and Waegeman [7] further emphasize that uncertainty must be communicated carefully because different forms of uncertainty influence how much confidence decision-makers should place in model outputs.

These studies justify using stochastic modeling in decision-support systems. Nevertheless, much of the literature either develops forecasting theory or reviews simulation methods without presenting a compact empirical procedure that converts operational data into state transitions, simulated paths, and managerial

risk indicators. This study therefore treats stochastic modeling not only as a mathematical exercise but also as a decision-support workflow.

Traffic forecasting has long been recognized as a central requirement for intelligent transportation operations. Vlahogianni et al. [27] show that short-term traffic forecasting is essential for traffic control and traveler information, while Lv et al. [14] and Ma et al. [15] demonstrate that deep learning and recurrent models can capture nonlinear temporal behavior in traffic data. Polson and Sokolov [20] extend this argument by showing that deep architectures can model sharp traffic transitions during special events.

Recent traffic studies increasingly account for spatial and temporal dependencies. Ermagun and Levinson [5] review spatiotemporal traffic forecasting and show that spatial dependence can improve prediction under congested conditions. Yu et al. [28] propose a spatio-temporal graph convolutional model, and Li et al. [12] introduce a diffusion convolutional recurrent model for traffic forecasting. These models are powerful, but they often require complex architectures and may produce less transparent outputs for operational decision-makers.

Probabilistic and simulation-based traffic studies more closely align with the present research. Sun et al. [25] use a Bayesian network logic for traffic-flow forecasting, Ran et al. [22] use attention-based recurrent learning for travel-time prediction, and Luan et al. [13] apply copula-based Monte Carlo simulation for probabilistic traffic prediction. Peng et al. [19] also use simulation under stochastic transport conditions. These studies show that uncertainty-aware traffic modeling matters, but they do not offer a simple cross-domain Markov-simulation framework applicable to both transportation and maintenance decision support.

Condition-based and predictive maintenance research provides a second domain where stochastic decision support is important. Alaswad and Xiang [2] review condition-based maintenance optimization models for stochastically deteriorating systems and show that maintenance decisions often depend on probabilistic deterioration. Susto et al. [26] present a multiple-classifier approach for predictive maintenance, while Kanawaday and Sane [8] explore time-series sensor modeling for industrial machines. These studies confirm that maintenance data can support operational decisions when deterioration or failure risk is modeled explicitly.

Broader predictive-maintenance reviews also emphasize the need for interpretable and deployable decision procedures. Lei et al. [10] review machinery prognostics from data acquisition to remaining-useful-life prediction, Carvalho et al. [3] review machine-learning methods in predictive maintenance, and Zonta et al. [29] synthesize predictive maintenance in Industry 4.0. Çınar et al. [4], Li et al. [11], Quatrini et al. [21], and Ran et al. [23] further show that maintenance systems require not only prediction but also operational integration, monitoring, risk communication, and decision timing.

More recent maintenance literature highlights deployment challenges. Achouch et al. [1] explain that predictive-maintenance systems are limited when sensor confidence, interpretability, or implementation readiness is weak. Morato et al. [17] and Papakonstantinou and Shinozuka [18] show that Markov decision processes and partially observable Markov models can support inspection and maintenance planning under uncertainty. These studies are mathematically strong, but they are often complex and domain-specific. The present study adopts a simpler empirical framework that keeps the stochastic state-transition logic while producing outputs that are easier to reproduce and explain.

The synthesis table summarizes the studies used in the literature review. It shows that traffic forecasting and predictive maintenance have advanced considerably, but researchers often study the two domains separately. Traffic studies emphasize prediction accuracy and spatiotemporal learning, while maintenance studies emphasize deterioration, failure risk, and maintenance policy. This paper addresses the gap of a compact empirical framework that represents random processes as state transitions, simulates future risk, and produces interpretable decision-support indicators across multiple operational domains.

Table 1. Literature-review synthesis of studies on stochastic modeling, traffic systems, and predictive maintenance.

Ref.	Focus	Method/design	Main contribution	Remaining gap for this study
[1]	Predictive maintenance challenges	Review of models and tools	Explained data-confidence and deployment challenges in maintenance systems.	Requires simple reproducible decision frameworks.
[2]	Condition-based maintenance	Review of optimization models	Showed how stochastic deterioration models support inspection and maintenance policy design.	Needs data-driven empirical validation with public datasets.
[3]	Machine learning in PdM	Systematic literature review	Reviewed machine-learning methods applied to predictive maintenance.	Highlights method diversity but not a unified stochastic DSS.
[4]	ML-based PdM	Review of smart manufacturing applications	Classified ML categories, data types, and maintenance applications.	Needs transparent uncertainty reporting for operational use.
[5]	Spatiotemporal traffic forecasting	Systematic review	Synthesized the role of spatial dependence in traffic prediction.	Does not provide a compact stochastic DSS template.
[6]	Probabilistic forecasting	Conceptual and statistical review	Clarified the need to evaluate full predictive distributions rather than point forecasts only.	Provides calibration logic but not an operational DSS workflow.
[7]	Uncertainty in predictive models	Conceptual review	Distinguished uncertainty sources and emphasized reliable probability communication.	Requires applied frameworks that translate uncertainty into action.
[8]	Industrial predictive maintenance	ARIMA/time-series sensor modeling	Explored forecasting sensor series for maintenance decision-making.	Limited probabilistic state-risk comparison.
[9]	Monte Carlo simulation	Methodological review	Showed simulation's central role in estimation, optimization, reliability, and operations research.	Does not integrate state-transition learning with decision outputs.
[10]	Machinery prognostics	Systematic review	Mapped the prognostics pipeline from data acquisition to remaining-life estimation.	More prognostics-oriented than decision-support oriented.
[11]	Condition-based maintenance	Prognostic and operational review	Connected prognostic information with production and maintenance decisions.	Does not empirically test a compact Markov-simulation framework.
[12]	Data-driven traffic forecasting	Diffusion convolutional recurrent network	Represented traffic movement through diffusion and recurrent temporal learning.	Requires network topology and extensive training resources.
[13]	Probabilistic traffic prediction	Copula-based Monte Carlo simulation	Modeled travel-time volatility and produced probabilistic traffic information.	Focuses on travel-time distribution rather than general DSS state transitions.
[14]	Traffic flow prediction	Deep learning model	Used stacked autoencoders to capture nonlinear traffic-flow patterns.	Strong point prediction, but limited interpretability for transition-based DSS.
[15]	Traffic speed prediction	LSTM model	Showed that recurrent models can represent temporal dependence in traffic sensors.	Model complexity reduces simple decision interpretability.

Ref.	Focus	Method/design	Main contribution	Remaining gap for this study
[17]	Inspection and maintenance planning	Dynamic Bayesian networks and POMDPs	Demonstrated stochastic decision models for deteriorating components.	Advanced stochastic control may be too complex for many applied DSS settings.
[18]	Structural inspection policies	Dynamic programming and Markov processes	Showed how Markov decision processes support sequential inspection planning.	The method is strong but domain-specific and computationally demanding.
[19]	Stochastic route planning	Monte Carlo and optimization	Used simulation for uncertain public transport route planning.	Application-specific and not transferable to maintenance systems.
[20]	Short-term traffic flow	Regularized deep learning	Captured sharp nonlinear changes in traffic flow during special events.	Limited focus on risk-state probabilities.
[21]	Condition-based maintenance	Bibliometric and literature review	Mapped major CBM clusters and research directions.	Limited empirical model demonstration.
[22]	Travel-time prediction	LSTM with attention	Improved temporal learning for travel-time prediction tasks.	Does not directly construct Markov risk states.
[23]	Predictive maintenance systems	Survey of systems, purposes, and methods	Reviewed architectures, optimization objectives, and predictive approaches.	General survey; lacks case-based state-transition analysis.
[25]	Traffic-flow forecasting	Spatio-temporal Bayesian network	Used probabilistic graphical modeling for traffic flow under uncertainty.	Probabilistic but not linked to Monte Carlo decision scenarios.
[26]	Predictive maintenance	Multiple classifier framework	Linked health factors to failure-risk classification for maintenance decisions.	Focuses on classification, not Markov transition interpretation.
[27]	Short-term traffic forecasting	Review of forecasting objectives and methods	Identified traffic forecasting as a core component of intelligent transport operations.	Most reviewed methods emphasize accuracy rather than state-risk decisions.
[28]	Traffic forecasting	Spatio-temporal graph convolution	Modeled traffic as graph-structured temporal data.	High-performing but less transparent for managers.
[29]	PdM in Industry 4.0	Systematic literature review	Identified sensor-based monitoring and predictive analytics as central maintenance enablers.	Implementation and decision integration remain open.

The reviewed studies support the need for an integrated stochastic decision-support framework that is simpler than high-dimensional black-box forecasting models but richer than deterministic summaries.

III. RESEARCH GAP

The reviewed literature reveals four gaps. First, many stochastic and probabilistic forecasting studies explain uncertainty well but do not always show how uncertainty becomes an operational decision output. Second, many traffic-flow studies prioritize prediction accuracy but give less attention to interpretable state transitions and scenario-based decision support. Third, many predictive-maintenance studies report failure prediction or review optimization policies but do not always connect state movement, probability, and simulation in a reproducible workflow. Fourth, few applied studies demonstrate the same stochastic decision-support framework across two different operational domains. This study addresses these gaps by combining traffic-flow and predictive-maintenance data in a unified Markov and Monte Carlo framework.

IV. STUDY CONTRIBUTION

This study makes three main contributions. First, it presents a practical stochastic-process modeling framework that integrates probability-distribution fitting, Markov state-transition modeling, and Monte Carlo simulation into a single decision-support procedure. Second, it demonstrates the framework using two empirical cases from different operational settings: traffic-flow congestion and machine-failure risk. Third, it reports outputs directly useful for decision support, including transition probabilities, the probability of entering a critical state, expected future vehicle counts, failure-risk probability, and validation metrics.

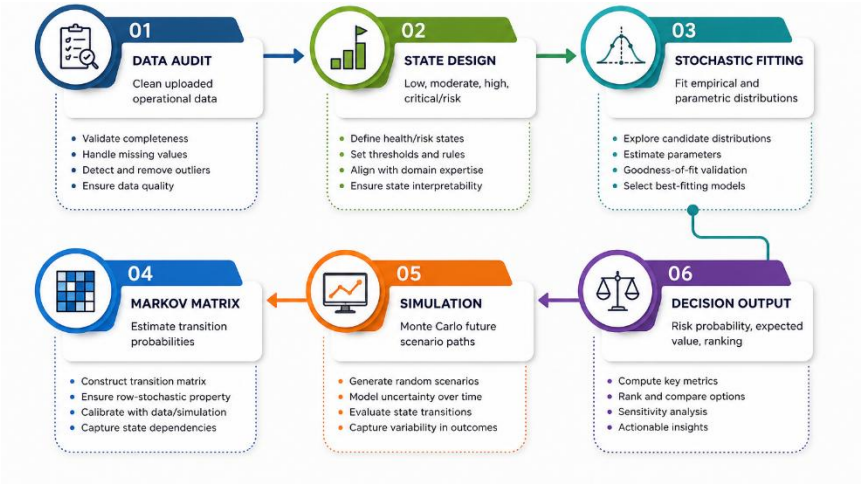

FIGURE 1. Proposed stochastic decision-support framework.

The figure shows how raw operational data are transformed into stochastic states, fitted distributions, transition matrices, simulated futures, and decision-support indicators.

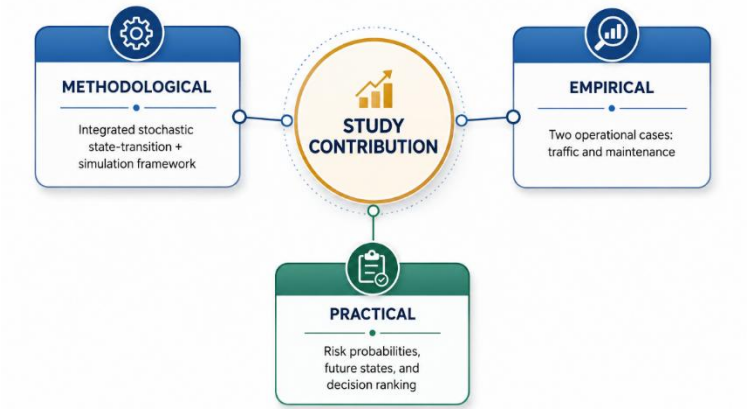

FIGURE 2. Study contribution map.

The figure summarizes the study's methodological, empirical, and practical contributions in a compact decision-support structure.

V. PROPOSED STOCHASTIC DECISION-SUPPORT FRAMEWORK

The proposed framework treats each operational system as a discrete-time stochastic process. Let the observed system state at time t be represented as:

$$X_t \in \mathcal{S} = \{s_1, s_2, \dots, s_k\}$$

Equation (1) defines the state space for the random process. In this study, the traffic case uses four congestion states, while the predictive-maintenance case uses four maintenance-risk states.

The transition probability from state s_i to state s_j is defined as:

$$p_{ij} = \Pr(X_{t+1} = s_j \mid X_t = s_i)$$

Equation (2) measures the probability that a system moves from one state to another in the next step.

The full transition matrix is given by:

$$\mathrm{P} = [p_{ij}], \quad \sum_{j=1}^{k} p_{ij} = 1, \quad p_{ij} \geq 0$$

Equation (3) ensures that each row of the Markov transition matrix forms a valid probability distribution.

For a state distribution vector $\boldsymbol{\pi}_t$, the expected future distribution after h steps is:

$$\pi_{t+h} = \pi_t \mathrm{P}^h$$

Equation (4) describes how future state probabilities are estimated over a decision horizon.

Monte Carlo simulation estimates future decision quantities by repeatedly sampling stochastic paths:

$$\hat{G} = \frac{1}{N} \sum_{r=1}^{N} g\left(X_{t+h}^{(r)}\right)$$

Equation (5) defines the Monte Carlo estimator for a decision quantity such as critical-state probability or expected future vehicle count.

The distribution-fitting stage uses the Akaike information criterion:

$$\mathrm{AIC} = 2k - 2\log(\hat{L})$$

Equation (6) compares candidate probability distributions by balancing goodness of fit and model complexity.

The multiclass Brier score for probabilistic state prediction is:

$$\mathrm{Brier} = \frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{k} (\hat{p}_{ij} - y_{ij})^2$$

Equation (7) evaluates how close predicted state probabilities are to the observed state outcomes.

Finally, a risk-sensitive decision indicator can be expressed as:

$$R_h = \Pr(X_{t+h} \in \mathcal{H}) \times E(Y_{t+h} \mid X_{t+h} \in \mathcal{H})$$

Equation (8) links the probability of high-risk states with the expected outcome severity inside those states.

VI. METHODOLOGY

1. RESEARCH DESIGN

The study uses an applied quantitative design based on two empirical cases. The main case uses hourly traffic-flow observations to model congestion-state transitions. The secondary case uses machine-operation observations to validate whether the same stochastic framework can support maintenance-related risk decisions. The analysis follows six stages: data audit, feature engineering, state construction, distribution fitting, Markov transition estimation, Monte Carlo simulation, and validation.

2. DATA SOURCES AND DATASET AUDIT

The traffic-flow dataset contains hourly observations from four road junctions and includes DateTime, junction, vehicle count, and an identifier [24]. The predictive-maintenance dataset contains operating variables from a synthetic milling process and includes operating conditions, product type, and machine-failure labels [16]. The uploaded Online Retail II files were inspected but excluded because they would shift

the manuscript toward customer-purchase behavior rather than operational random-process decision support.

Table 2. Dataset audit for the two empirical cases.

Case	Metric	Value
Traffic flow	Observations	48,120
Traffic flow	Original variables	4
Traffic flow	Engineered variables used	Hour, day of week, month, weekend, state, next state
Traffic flow	Junctions	4
Traffic flow	Date range	2015-11-01 to 2017-06-30
Traffic flow	Missing values in raw data	0
Traffic flow	Duplicate rows in raw data	0
Traffic flow	Mean vehicles per hour	22.79
Traffic flow	Std. dev. vehicles	20.75
Traffic flow	Min-max vehicles	1-180
Predictive maintenance	Observations	10,000
Predictive maintenance	Original variables	14
Predictive maintenance	Engineered variables used	Failure probability, risk state, next risk state
Predictive maintenance	Missing values in raw data	0
Predictive maintenance	Duplicate rows in raw data	0
Predictive maintenance	Machine failures	339 (3.39%)

The dataset audit shows that the selected traffic and maintenance datasets are clean enough for stochastic modeling because both contain no missing values or duplicate rows in the inspected raw files.

3. TRAFFIC-FLOW CASE STUDY

We sorted the traffic data by junction and time. Four state categories were created separately for each junction using quartile thresholds. This state design avoids treating all junctions as if they had the same traffic scale. We classified vehicle counts at or below the first quartile as low; values between the first quartile and median as moderate; values between the median and third quartile as high; and values above the third quartile as critical. We then created a next-hour state within each junction to estimate Markov transition probabilities.

Table 3. Descriptive statistics by traffic junction.

Junction	n	Mean	Std. dev.	Min	Median	Max
1	14592	45.05	23.01	5	40.0	156
2	14592	14.25	7.40	1	13.0	48
3	14592	13.69	10.44	1	11.0	180
4	4344	7.25	3.52	1	7.0	36

The table shows that Junction 1 has the highest average vehicle count, while Junction 4 has the lowest average volume and a shorter observation record.

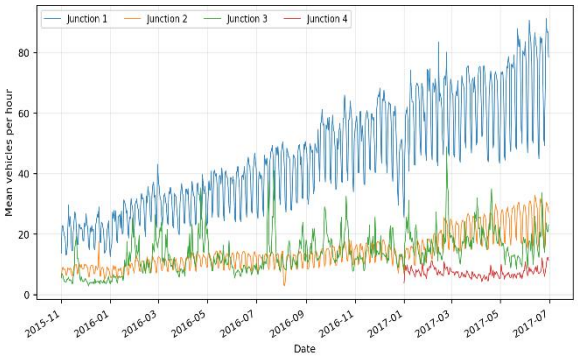

FIGURE 3. Daily traffic-flow pattern by junction.

The figure shows time-dependent variation in daily traffic volumes and supports modeling traffic as a stochastic process rather than a fixed deterministic sequence.

4. PREDICTIVE-MAINTENANCE VALIDATION CASE

The predictive-maintenance case uses operating variables including air temperature, process temperature, rotational speed, torque, tool wear, product type, and the machine-failure label. We trained a probabilistic classifier to estimate failure probability, then converted the predicted probability into four risk states: normal, watch, high-risk, and failure. We assigned the failure state directly to observations with a machine-failure label, while assigning non-failure observations to normal, watch, or high-risk states using probability thresholds derived from the non-failure probability distribution.

5. ANALYTICAL PROCEDURE

The traffic case used a time-aware train-test split by junction. The first 80% of each junction sequence was used for transition estimation, and the remaining 20% was used for one-step state validation. Traffic distribution fitting compared normal, lognormal, gamma, and negative-binomial models using AIC. The maintenance case used a stratified split for probabilistic validation. A random-forest classifier estimated failure probability, and we used a probability threshold of 0.30 to identify failure-risk predictions for classification metrics. Monte Carlo simulation used the estimated transition matrices to generate future state paths.

6. VALIDITY, RELIABILITY, AND ETHICS

Validity was supported by using two different empirical cases, by separating traffic model estimation from traffic validation, and by reporting probabilistic metrics rather than accuracy alone. We supported reliability through transparent state definitions, fixed simulation seeds, and reproducible transition-matrix estimation. Because the data are public operational datasets and contain no direct human-identifying information, the study does not require participant-level ethical review. The analysis remains limited by the original data structure and by the assumption that estimated transition probabilities remain stable over the simulated horizon.

VII. RESULTS

1. DISTRIBUTION FITTING FOR TRAFFIC VEHICLE COUNTS

Table 4. Distribution fitting for traffic vehicle counts.

Distribution	Parameters	k	AIC
Lognormal	shape=0.849, scale=15.997	2	387646.400
Gamma	alpha=1.558, scale=14.626	2	392081.800

Distribution	Parameters	k	AIC
Negative binomial	r=1.274, p=0.053	2	395121.300
Normal	mu=22.791, sigma=20.750	2	428414.200

The lognormal distribution has the lowest AIC among the tested candidates, indicating that traffic vehicle counts are positively skewed and better represented by a non-normal distribution.

2. MARKOV STATE TRANSITIONS IN TRAFFIC FLOW

Table 5. Full-sample traffic Markov transition probabilities.

State	Low	Moderate	High	Critical
Low	0.760	0.209	0.025	0.006
Moderate	0.243	0.532	0.200	0.025
High	0.031	0.205	0.592	0.172
Critical	0.005	0.024	0.177	0.794

The matrix shows strong within-state persistence, especially for low and critical states, while moderate and high states show more frequent movement to neighboring states.

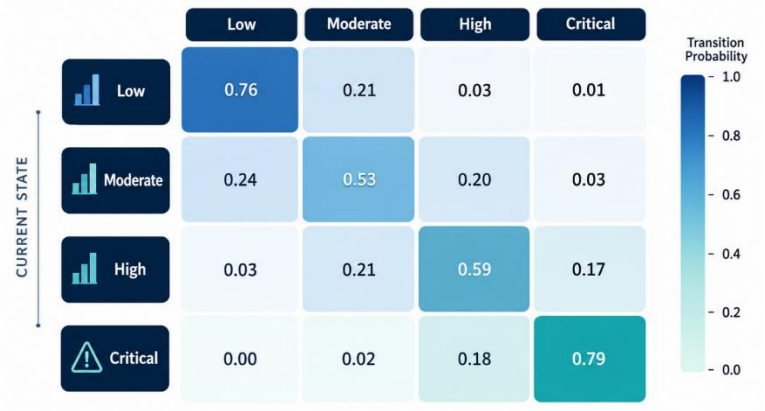

FIGURE 4. Traffic Markov transition heatmap.

The heatmap shows that critical congestion is highly persistent, which matters for short-horizon traffic-control decisions.

3. MONTE CARLO SIMULATION FOR TRAFFIC DECISION SUPPORT

Table 6. Monte Carlo traffic-risk simulation over a 24-hour horizon.

Junction	Initial state	Initial vehicles	P(any critical within 24h)	Mean P(critical per hour)	P critical at 24h	Expected vehicles at 24h
1	Critical	78	0.953	0.352	0.247	44.805
2	Critical	27	0.954	0.358	0.246	14.436
3	Critical	39	0.952	0.350	0.228	13.653

Junction	Initial state	Initial vehicles	P(any critical within 24h)	Mean P(critical per hour)	P critical at 24h	Expected vehicles at 24h
4	Critical	12	0.950	0.356	0.231	7.568

The simulation shows that when a junction starts in a critical state, the probability of experiencing at least one critical hour over the next 24 hours exceeds 0.95 for all junctions.

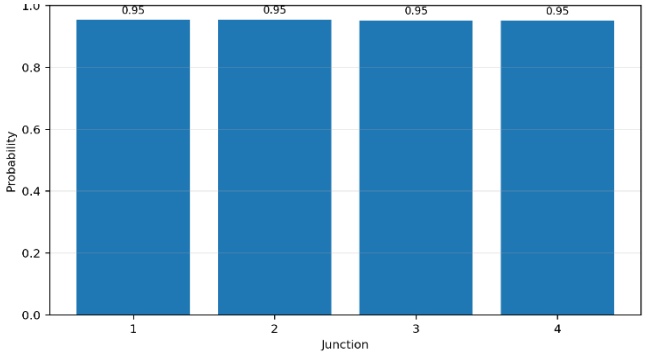

FIGURE 5. Monte Carlo simulation of critical traffic risk.

The figure translates transition probabilities into direct risk indicators that can support congestion monitoring and resource allocation.

4. PREDICTIVE-MAINTENANCE STATE TRANSITIONS

Table 7. Predictive-maintenance Markov transition probabilities.

RiskState	Normal	Watch	High-risk	Failure
Normal	0.583	0.348	0.054	0.016
Watch	0.435	0.436	0.093	0.036
High-risk	0.282	0.319	0.300	0.099
Failure	0.301	0.445	0.168	0.086

The maintenance transition matrix shows that high-risk states have a higher direct transition probability to failure than normal and watch states.



FIGURE 6. Predictive-maintenance transition heatmap.

The heatmap shows how risk states escalate toward failure and supports maintenance prioritization before failure occurs.

5. MAINTENANCE-RISK SIMULATION AND MODEL VALIDATION

Table 8. Maintenance-risk Monte Carlo simulation.

Initial risk state	Horizon	P(any failure)
Watch	20 process steps	0.477
High-risk	20 process steps	0.536
Watch	100 process steps	0.960
High-risk	100 process steps	0.965

The simulation shows that the probability of at least one failure event rises sharply as the horizon increases from 20 to 100 process steps.

Table 9. Model validation results.

Metric	Value
Traffic Markov one-step state accuracy	0.720
Traffic persistence baseline accuracy	0.720
Traffic multiclass Brier score	0.430
Traffic multiclass log loss	0.785
AI4I probabilistic validation AUC	0.976
AI4I Brier score	0.014
AI4I accuracy	0.983
AI4I precision	0.734
AI4I recall	0.784
AI4I F1 score	0.758

The validation results show that the traffic Markov model is interpretable and calibrated for state movement, while the maintenance model provides strong probabilistic discrimination for failure-risk estimation.

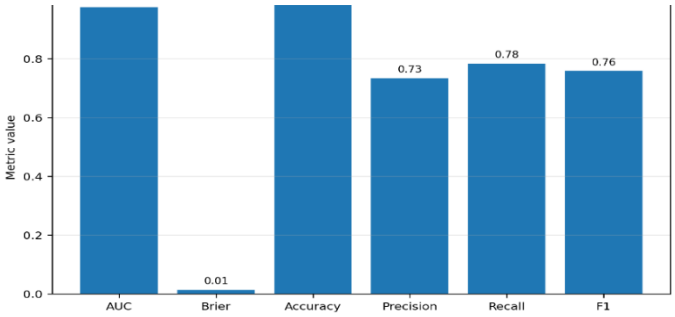

FIGURE 7. Predictive-maintenance validation metrics.

The figure shows strong failure discrimination, although precision and F1 indicate that rare-failure prediction remains more difficult than overall accuracy suggests.

VIII. DISCUSSION

The empirical findings support the value of stochastic modeling for decision support under uncertainty. The traffic case demonstrates that hourly vehicle movement can be represented through discrete congestion states and transition probabilities. The high persistence of low and critical states matters in practice. A low state tends to remain stable, while a critical state is not merely a temporary spike; it often continues or returns within the short decision horizon. This means that traffic managers should not rely only on the current vehicle count. They should also examine the probability that the state will remain critical or return to critical conditions after short-term relief.

The distribution-fitting results also strengthen the stochastic interpretation of traffic flow. The lognormal distribution outperforms the normal, gamma, and negative-binomial alternatives according to AIC. This finding is expected because vehicle counts are positive and skewed. From a decision-support perspective, this implies that average-based traffic indicators are incomplete. Systems with skewed outcomes require probability-based thresholds and scenario estimates because extreme congestion events carry more operational importance than ordinary fluctuations.

The one-step traffic validation result requires careful interpretation. The Markov model achieves an accuracy of 0.720, matching the persistence baseline. This does not mean that the Markov model is unnecessary. Instead, it means traffic states are highly persistent and that a simple current-state rule is already hard to outperform with only a first-order transition matrix. The Markov model's value is not only classification accuracy; it also provides a full probability distribution over future states. This lets decision-makers evaluate risk, not just predict the most likely next state.

The Monte Carlo simulation results provide the clearest decision-support output in the traffic case. When the current state is critical, the probability of encountering at least one critical hour within the next 24 hours exceeds 0.95 across all junctions. This result is actionable because it converts a current condition into a future risk statement. A control center could use this information to prioritize monitoring, signal control, staffing, traveler warnings, or additional data collection at high-risk junctions.

The predictive-maintenance case confirms that the proposed framework is transferable beyond traffic. The failure-risk model achieves high discrimination (AUC = 0.976), and the state-transition matrix shows meaningful escalation from high-risk to failure states. However, the precision and F1 score show that rare-failure prediction remains challenging. This is a realistic pattern in maintenance datasets where failure events are much less frequent than normal operating observations. The stochastic framework is therefore useful because it reports probabilities and risk-state movement rather than a single binary prediction.

The maintenance simulation is also important. The probability of at least one failure event grows from approximately 0.477 to 0.960 for watch-state observations as the horizon expands from 20 to 100 process steps, and from approximately 0.536 to 0.965 for high-risk observations. This does not imply that every individual machine will fail with certainty, because the simulation uses empirical transition probabilities and repeated scenario paths. However, it does show that the time horizon strongly affects maintenance risk. A risk state that appears acceptable over a short horizon may become unacceptable over a longer planning horizon.

Across both cases, the study shows that stochastic decision support should combine three outputs: present-state classification, transition probability, and simulated future risk. A classification label alone may be too rigid. A transition matrix alone may be too abstract. A simulation alone may be difficult to interpret without clear state definitions. When combined, these components provide a practical evidence chain from data to decision.

The proposed framework is intentionally simpler than deep spatiotemporal architectures or advanced stochastic-control models. This simplicity is a strength for applied decision-support settings where transparency, reproducibility, and communication matter. Managers can inspect the state definitions, transition probabilities, and simulation outputs. Analysts can reproduce the results without specialized

graph structures or black-box optimization. The framework therefore supports practical decision-making while preserving a clear mathematical foundation.

IX. THEORETICAL CONTRIBUTIONS

Theoretically, the study contributes to stochastic decision-support research by integrating probability distributions, Markov chains, and Monte Carlo simulation into a unified applied framework. It shows that random processes can be represented not only as numerical time series but also as interpretable state systems. This helps connect stochastic-process theory with operational decision support.

The study also shows that accuracy is not the only relevant criterion for stochastic models. In the traffic case, the Markov model does not outperform the persistence baseline in one-step state accuracy, but it adds probabilistic information that the baseline cannot provide. This distinction matters for theory because decision support under uncertainty depends on probability distributions, risk states, and scenario paths, not just point predictions.

Finally, the two-case design extends the framework's generality. Traffic congestion and machine failure are different operational problems, yet both can be represented through stochastic states and transition matrices. This supports the theoretical argument that stochastic decision-support systems can be designed around state evolution rather than domain-specific prediction alone.

X. XI. LIMITATIONS AND FUTURE RESEARCH

The study has several limitations. First, the traffic dataset excludes weather, road incidents, public events, and signal-control variables, which may influence congestion-state transitions. Second, the Markov model assumes first-order transitions, meaning the next state depends only on the current state. Future studies could test higher-order Markov models, hidden Markov models, or semi-Markov models. Third, the maintenance dataset is synthetic, so future work should validate the framework using additional industrial sensor datasets with real failure histories. Fourth, the Monte Carlo simulation assumes that estimated transition probabilities remain stable across the forecast horizon.

Future research can extend the framework in four ways. First, add exogenous predictors such as weather, calendar events, sensor condition, and operating load to create conditional transition matrices. Second, include cost functions to rank decisions by expected loss rather than probability alone. Third, introduce dynamic updating so transition probabilities recalibrate as new observations arrive. Fourth, the framework can be tested in healthcare, finance, logistics, and energy systems where random processes and risk-sensitive decisions are central.

XII. CONCLUSION

This study developed and evaluated a stochastic process modeling framework for decision support under uncertainty. Using traffic-flow data and predictive-maintenance data, the study showed how random operational behavior can be transformed into stochastic states, transition probabilities, simulated future scenarios, and interpretable risk indicators. The traffic case showed that vehicle counts are skewed, critical congestion states are highly persistent, and Monte Carlo simulation can estimate the probability of future critical conditions. The maintenance case demonstrated that probabilistic failure modeling and Markov risk states can support maintenance prioritization across different planning horizons.

The main conclusion is that stochastic modeling strengthens decision-support systems by moving analysis beyond deterministic prediction. A useful decision-support system should not only state what is most likely to happen; it should also show how likely alternative outcomes are, how risks evolve over time, and how future scenarios may affect operational decisions. The proposed framework offers a practical, reproducible way to achieve this transformation. Although the framework should be extended with richer predictors and additional real-world validation, the results support its value for uncertainty-aware decision support in dynamic operational systems.

Author Contributions

The author conducted the conceptualization, methodology, data analysis, investigation, writing, review, editing, and final approval of the manuscript.

Funding

This research received no external funding.

Data Availability

The dataset will be available from the author upon reasonable request.

Conflicts of Interest

The author declares no conflict of interest.

REFERENCES

[1]. Achouch, M., Dimitrova, M., Ziane, K., Sattarpanah Karganroudi, S., Dhouib, R., Ibrahim, H., & Adda, M. (**2022**). On predictive maintenance in Industry 4.0: Overview, models, and challenges. *Applied Sciences, 12*(16), 8081. https://doi.org/10.3390/app12168081

[2]. Alaswad, S., & Xiang, Y. (**2017**). A review on condition-based maintenance optimization models for stochastically deteriorating systems. *Reliability Engineering & System Safety, 157*, 54–63. https://doi.org/10.1016/j.ress.2016.08.009

[3]. Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Basto, J. P., & Alcalá, S. G. S. (**2019**). A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering, 137*, 106024. https://doi.org/10.1016/j.cie.2019.106024

[4]. Çınar, Z. M., Nuhu, A. A., Zeeshan, Q., Korhan, O., Asmael, M., & Safaei, B. (**2020**). Machine learning in predictive maintenance towards sustainable smart manufacturing in Industry 4.0. *Sustainability, 12*(19), 8211. https://doi.org/10.3390/su12198211

[5]. Ermagun, A., & Levinson, D. (**2018**). Spatiotemporal traffic forecasting: Review and proposed directions. *Transport Reviews, 38*(6), 786–814. https://doi.org/10.1080/01441647.2018.1442887

[6]. Gneiting, T., & Katzfuss, M. (**2014**). Probabilistic forecasting. *Annual Review of Statistics and Its Application, 1*, 125–151. https://doi.org/10.1146/annurev-statistics-062713-085831

[7]. Hüllermeier, E., & Waegeman, W. (**2021**). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning, 110*, 457–506. https://doi.org/10.1007/s10994-021-05946-3

[8]. Kanawaday, A., & Sane, A. (**2017**). Machine learning for predictive maintenance of industrial machines using IoT sensor data. *2017 8th IEEE International Conference on Software Engineering and Service Science*, 87–90. https://doi.org/10.1109/ICSESS.2017.8342870

[9]. Kroese, D. P., Brereton, T., Taimre, T., & Botev, Z. I. (**2014**). Why the Monte Carlo method matters today. *WIREs Computational Statistics, 6*(6), 386–392. https://doi.org/10.1002/wics.1314

[10]. Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (**2018**). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing, 104*, 799–834. https://doi.org/10.1016/j.ymssp.2017.11.016

[11]. Li, Y., Peng, S., Li, Y., & Jiang, W. (**2020**). A review of condition-based maintenance: Its prognostic and operational aspects. *Frontiers of Engineering Management, 7*, 323–334. https://doi.org/10.1007/s42524-020-0121-5

[12]. Li, Y., Yu, R., Shahabi, C., & Liu, Y. (**2018**). Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *International Conference on Learning Representations*. https://doi.org/10.48550/arXiv.1707.01926

[13]. Luan, S., Chen, X., Su, Y., Dong, Z., & Ma, X. (**2022**). Modeling travel time volatility using copula-based Monte Carlo simulation method for probabilistic traffic prediction. *Transportmetrica A: Transport Science, 18*(3), 382–408. https://doi.org/10.1080/23249935.2019.1692959

[14]. Lv, Y., Duan, Y., Kang, W., Li, Z., & Wang, F.-Y. (**2015**). Traffic flow prediction with big data: A deep learning approach. *IEEE Transactions on Intelligent Transportation Systems, 16*(2), 865–873. https://doi.org/10.1109/TITS.2014.2345663

[15]. Ma, X., Tao, Z., Wang, Y., Yu, H., & Wang, Y. (**2015**). Long short-term memory neural network for traffic speed prediction using remote microwave sensor data. *Transportation Research Part C: Emerging Technologies, 54*, 187–197. https://doi.org/10.1016/j.trc.2015.03.014

[16]. Matzka, S. (**2022**). Predictive maintenance dataset AI4I 2020. *Kaggle Dataset*. https://www.kaggle.com/datasets/stephanmatzka/predictive-maintenance-dataset-ai4i-2020

[17]. Morato, P. G., Papakonstantinou, K. G., Andriotis, C. P., Nielsen, J. S., & Rigo, P. (**2022**). Optimal inspection and maintenance planning for deteriorating structural components through dynamic Bayesian networks and Markov decision processes. *Structural Safety*, *94*, 102140. https://doi.org/10.1016/j.strusafe.2021.102140

[18]. Papakonstantinou, K. G., & Shinozuka, M. (**2014**). Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation. *Reliability Engineering & System Safety*, *130*, 214–224. https://doi.org/10.1016/j.ress.2014.04.006

[19]. Peng, Y., Meng, Q., Li, Z., & Wang, W. (**2021**). Passengers' route planning in stochastic common-lines public transport networks. *Archives of Transport*, *59*(3), 7–22. https://doi.org/10.5604/01.3001.0015.2643

[20]. Polson, N. G., & Sokolov, V. O. (**2017**). Deep learning for short-term traffic flow prediction. *Transportation Research Part C: Emerging Technologies*, *79*, 1–17. https://doi.org/10.1016/j.trc.2017.02.024

[21]. Quatrini, E., Costantino, F., Di Gravio, G., & Patriarca, R. (**2020**). Condition-based maintenance: An extensive literature review. *Machines*, *8*(2), 31. https://doi.org/10.3390/machines8020031

[22]. Ran, X., Shan, Z., Fang, Y., & Lin, C. (**2019**). An LSTM-based method with attention mechanism for travel time prediction. *Sensors*, *19*(4), 861. https://doi.org/10.3390/s19040861

[23]. Ran, Y., Zhou, X., Lin, P., Wen, Y., & Deng, R. (**2019**). A survey of predictive maintenance: Systems, purposes and approaches. *arXiv*. https://doi.org/10.48550/arXiv.1912.07383

[24]. Soriano, F. (**2021**). Traffic prediction dataset. *Kaggle Dataset*. https://www.kaggle.com/datasets/fedesoriano/traffic-prediction-dataset

[25]. Sun, S., Zhang, C., & Zhang, Y. (**2017**). Traffic flow forecasting using a spatio-temporal Bayesian network predictor. *arXiv*. https://doi.org/10.48550/arXiv.1712.08883

[26]. Susto, G. A., Schirru, A., Pampuri, S., McLoone, S., & Beghi, A. (**2015**). Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, *11*(3), 812–820. https://doi.org/10.1109/TII.2014.2349359

[27]. Vlahogianni, E. I., Karlaftis, M. G., & Golias, J. C. (**2014**). Short-term traffic forecasting: Where we are and where we are going. *Transportation Research Part C: Emerging Technologies*, *43*, 3–19. https://doi.org/10.1016/j.trc.2014.01.005

[28]. Yu, B., Yin, H., & Zhu, Z. (**2018**). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *Proceedings of IJCAI*, 3634–3640. https://doi.org/10.24963/ijcai.2018/505

[29]. Zonta, T., da Costa, C. A., da Rosa Righi, R., de Lima, M. J., da Trindade, E. S., & Li, G. P. (**2020**). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering*, *150*, 106889. https://doi.org/10.1016/j.cie.2020.106889